

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	2198	newsgroup or usenet	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:38
S2	314	(weighted adj graph) or (spectral adj clustering)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:51
S3	3	S1 and S2	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:19
S5	2	S1 with S2	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:19
S6	99308	newsgroup or usenet or forum or (special adj interest adj group) or sig	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:21
S7	8	S6 and S2	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:23
S8	2	S6 same S2	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:22
S9	2	S6 and (cross-post\$3 or crosspost\$3) and graph	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:44
S10	1	"5796393".PN.	USPAT; USOCR	OR	OFF	2006/01/04 15:45

S11	1	"6215495".PN.	USPAT; USOCR	OR	OFF	2006/01/04 15:46
S12	6	("5796393" "6215495" "6266805" "6289299" "6295514").PN. OR ("6594673").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/01/04 15:46
S13	280	(weighted adj graph)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:51
S14	280	(weighted adj graph)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:37
S15	37	S14 and "707".clas.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 15:54
S16	2198	newsgroup or usenet	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 16:23
S17	314	(weighted adj graph) or (spectral adj clustering)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:32
S18	2	S16 same S17	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 17:16
S19	5	S16 and (cross-post\$3 or cross adj post\$3 or crosspost\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/04 17:17

S20	2206	newsgroup or usenet	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:42
S21	3	S20 and (weighted adj graph)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:37
S22	4	(newsgroup or usenet) and cross-post\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:38
S23	284	weighted adj graph	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:43
S24	60	spectral adj cluster\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/05 11:44
S25	180685	(newsgroup\$1 board adj message\$1 bbs bulletin)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:06
S26	959163	(graph\$1 chart\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:03
S27	24380	S25 and S26	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:04

S28	777	S25 same S26	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:04
S29	5915188	(generat\$3 creat\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:05
S30	959163	(graph\$1 chart\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:04
S31	3146670	(cluster\$3 group\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:05
S32	6584716	(generat\$3 creat\$3 build\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:06
S33	192723	(newsgroup\$1 board adj message\$1 bbs bulletin forum)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:06
S34	14483	S33 and S31 and S32 and S30	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:08
S35	3	S33 with S31 with S32 with S30	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:16

S36	26	S33 same S31 same S32 same S30	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:10
S37	2	"5923846".pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:18
S38	2	"6336132".pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:18
S39	1	"6952700".pn.	USPAT	OR	OFF	2006/01/06 14:31
S40	164662	"707".clas. "709".clas. "345".clas. "711".clas.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:35
S41	318	(weighted adj graph) or (spectral adj clustering)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:35
S42	85	S41 and S40	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:35
S43	959163	(graph\$1 chart\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:35
S44	3146670	(cluster\$3 group\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:35

S45	6584716	(generat\$3 creat\$3 build\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:35
S46	192723	(newsgroup\$1 board adj message\$1 bbs bulletin forum)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:35
S47	14483	S46 and S44 and S45 and S43	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:35
S48	2462	S47 and S40	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:37
S49	26	S46 same S44 same S45 same S43	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/06 14:37
S50	13	S49 and S40	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/01/06 14:38

Lowercase "or" was ignored. Try "OR" to search for either of two terms. [\[details\]](#)
The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Scholar

Results 1 - 4 of 4 for **(newsgroup or usenet) and ("weighted graph")**. (0.05 seconds)

Tip: Try removing quotes from your search to get more results.

Exploring the community structure of newsgroups

C Borgs, J Chayes, M Mahdian, A Saberi - Proceedings of the 2004 ACM SIGKDD international conference ... , 2004 - portal.acm.org

... the cross-post graph as a **weighted graph** with vertices ... <http://research.microsoft.com/jchayes/Papers/usenet.html> ... to help users find the right **newsgroup** to post ...

[Cited by 2](#) - Web Search - research.microsoft.com - research.microsoft.com - portal.acm.org

An Intelligent Agent for High-Precision Text Filtering

A O'Riordan, H Sorensen - CIKM, 1995 - portal.acm.org

... A **weighted graph** representation is used for documents, and graph manipulation algorithms are used in the processing. 1 Introduction ...

[Cited by 8](#) - Web Search - cdserv4.inria.fr - portal.acm.org

Analyzing the Effectiveness and Applicability of Co-training

K Nigam, R Ghani - CIKM, 2000 - portal.acm.org

... Thus, the word 'career' from the rest **newsgroup** is a distinct feature from the ... When tokenizing this data, the **UseNet** headers (including the subject line) are ...

[Cited by 109](#) - Web Search - cs.wustl.edu - kamalnigam.com - www-2.cs.cmu.edu - [all 13 versions](#) »

On Scaling Up Balanced Clustering Algorithms

A Banerjee, J Ghosh - SDM, 2002 - lans.ece.utexas.edu

... a graph partitioning problem [22]. A **weighted graph** is constructed whose vertices are the data-points. An edge connecting any two ...

[Cited by 9](#) - [View as HTML](#) - Web Search - lans.ece.utexas.edu

Terms used newsgroup or usenet and weighted and graph

Found 14,608 of 169,166

Sort results by

 Save results to a Binder[Try an Advanced Search](#)

Display results

 Search Tips[Try this search in The ACM Guide](#) Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [Industry/government track posters: Exploring the community structure of newsgroups](#)

Christian Borgs, Jennifer Chayes, Mohammad Mahdian, Amin Saberi

 August 2004 **Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '04****Publisher:** ACM PressFull text available: [pdf\(3.11 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We propose to use the community structure of Usenet for organizing and retrieving the information stored in newsgroups. In particular, we study the network formed by cross-posts, messages that are posted to two or more newsgroups simultaneously. We present what is, to our knowledge, by far the most detailed data that has been collected on Usenet cross-postings. We analyze this network to show that it is a small-world network with significant clustering. We also present a spectral algorithm which ...

Keywords: clustering, spectral method, usenet**2** [Data mining: Mining newsgroups using networks arising from social behavior](#)

Rakesh Agrawal, Sridhar Rajagopalan, Ramakrishnan Srikant, Yirong Xu

 May 2003 **Proceedings of the 12th international conference on World Wide Web****Publisher:** ACM PressFull text available: [pdf\(299.31 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Recent advances in information retrieval over hyperlinked corpora have convincingly demonstrated that links carry less noisy information than text. We investigate the feasibility of applying link-based methods in new applications domains. The specific application we consider is to partition authors into opposite camps within a given topic in the context of newsgroups. A typical newsgroup posting consists of one or more quoted lines from another posting followed by the opinion of the author. This ...

Keywords: data mining, link analysis, newsgroup, social network, text mining, web mining**3** [Clustering: Bipartite graph partitioning and data clustering](#)

Hongyuan Zha, Xiaofeng He, Chris Ding, Horst Simon, Ming Gu

 October 2001 **Proceedings of the tenth international conference on Information and knowledge management****Publisher:** ACM PressFull text available: [pdf\(1.45 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Many data types arising from data mining applications can be modeled as bipartite graphs, examples include terms and documents in a text corpus, customers and purchasing items in market basket analysis and reviewers and movies in a movie

recommender system. In this paper, we propose a new data clustering method based on partitioning the underlying bipartite graph. The partition is constructed by minimizing a normalized sum of edge weights between *unmatched* pairs of vertices of the ...

Keywords: bipartite graph, correspondence analysis, document clustering, graph partitioning, singular value decomposition, spectral relaxation

4 GroupLens: an open architecture for collaborative filtering of netnews

 Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, John Riedl
October 1994 **Proceedings of the 1994 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available:  pdf(1.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Collaborative filters help people make choices based on the opinions of other people. GroupLens is a system for collaborative filtering of netnews, to help people find articles they will like in the huge stream of available articles. News reader clients display predicted scores and make it easy for users to rate articles after they read them. Rating servers, called Better Bit Bureaus, gather and disseminate the ratings. The rating servers predict scores based on the heuristic that people wh ...

Keywords: Usenet, collaborative filtering, electronic bulletin boards, information filtering, netnews, selective dissemination of information, social filtering, user model

5 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren
November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available:  pdf(4.21 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

6 The SIFT information dissemination system

 Tak W. Yan, Hector Garcia-Molina
December 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 4

Publisher: ACM Press

Full text available:  pdf(220.77 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Information dissemination is a powerful mechanism for finding information in wide-area environments. An information dissemination server accepts long-term user queries, collects new documents from information sources, matches the documents against the queries, and continuously updates the users with relevant information. This paper is a retrospective of the Stanford Information Filtering Service (SIFT), a system that as of April 1996 was processing over 40,000 worldwide subscriptions and ov ...

Keywords: Boolean queries, dissemination, filtering, indexing, vector space queries

7 Data mining: A matrix density based algorithm to hierarchically co-cluster documents and words

 Bhushan Mandhani, Sachindra Joshi, Krishna Kummamuru
May 2003 **Proceedings of the 12th international conference on World Wide Web**

Publisher: ACM Press

This paper proposes an algorithm to hierarchically cluster documents. Each cluster is actually a cluster of documents and an associated cluster of words, thus a document-word co-cluster. Note that, the vector model for documents creates the document-word matrix, of which every co-cluster is a submatrix. One would intuitively expect a submatrix made up of high values to be a good document cluster, with the corresponding word cluster containing its most distinctive features. Our algorithm looks to ...

8 Optimistic replication

 Yasushi Saito, Marc Shapiro

March 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 1

Publisher: ACM Press

Full text available:  pdf(656.72 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes popular. Optimistic replication deploys algorithms not seen in tradition ...

Keywords: Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques

9 An intelligent agent for high-precision text filtering

 Adrian O'Riordan, Humphrey Sorensen

December 1995 **Proceedings of the fourth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available:  pdf(636.79 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

10 Research session 5: data mining / transaction management: A divide-and-merge methodology for clustering

 David Cheng, Santosh Vempala, Ravi Kannan, Grant Wang

June 2005 **Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  pdf(791.76 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

We present a divide-and-merge methodology for clustering a set of objects that combines a top-down "divide" phase with a bottom-up "merge" phase. In contrast, previous algorithms either use top-down or bottom-up methods to construct a hierarchical clustering or produce a flat clustering using local search (e.g., k -means). Our divide phase produces a tree whose leaves are the elements of the set. For this phase, we use an efficient spectral algorithm. The merge phase quickly finds an optim ...

11 Customizing information capture and access

 Daniela Rus, Devika Subramanian

January 1997 **ACM Transactions on Information Systems (TOIS)**, Volume 15 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.26 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article presents a customizable architecture for software agents that capture and access information in large, heterogeneous, distributed electronic repositories. The key idea is to exploit underlying structure at various levels of granularity to build high-level indices with task-specific interpretations. Information agents construct such indices and are configured as a network of reusable modules called structure detectors and segmenters. We illustrate our architectu ...

Keywords: information gathering, software agents, table recognition

12 Communities: Flash forums and forumReader: navigating a new kind of large-scale online discussion

 Kushal Dave, Martin Wattenberg, Michael Muller

November 2004 **Proceedings of the 2004 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available:  [pdf\(513.95 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We describe a popular kind of large, topic-centered, transient discussion, which we term a *flash forum*. These occur in settings ranging from web-based bulletin boards to corporate intranets, and they display a conversational style distinct from Usenet and other online discussion. Notably, authorship is more diffuse, and threads are less deep and distinct. To help orient users and guide them to areas of interest within flash forums, we designed ForumReader, a tool combining data ...

Keywords: collaboration, large-scale conversations, mass interaction, persistent conversations, prototype, thumbnail interface, user interface, user study, visualization

13 Experiments in social data mining: The TopicShop system

 Brian Amento, Loren Terveen, Will Hill, Deborah Hix, Robert Schulman

March 2003 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 10 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(377.92 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Social data mining systems enable people to share opinions and benefit from each other's experience. They do this by mining and redistributing information from computational records of social activity such as Usenet messages, system usage history, citations, or hyperlinks. Some general questions for evaluating such systems are: (1) is the extracted information valuable? and (2) do interfaces based on the information improve user task performance? We report here on *TopicShop*, a syst ...

Keywords: Cocitation analysis, collaborative filtering, computer-supported cooperative work, information visualization, social filtering, social network analysis

14 Constructing, organizing, and visualizing collections of topically related Web

 resources

 Loren Terveen, Will Hill, Brian Amento

March 1999 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 6 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(303.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

For many purposes, the Web page is too small a unit of interaction and analysis. Web sites are structured multimedia documents consisting of many pages, and users often are interested in obtaining and evaluating entire collections of topically related sites. Once such a collection is obtained, users face the challenge of exploring, comprehending and organizing the items. We report four innovations that address these user needs: (1) we replaced the Web page with the Web site

Keywords: cocitation analysis, collaborative filtering, computer supported cooperative work, information visualization, social filtering, social network analysis

15 Research track papers: A probabilistic framework for semi-supervised clustering

 Sugato Basu, Mikhail Bilenko, Raymond J. Mooney

August 2004 **Proceedings of the tenth ACM SIGKDD international conference on**

Publisher: ACM Press

Full text available: [pdf\(187.51 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Unsupervised clustering can be significantly improved using supervision in the form of pairwise constraints, i.e., pairs of instances labeled as belonging to same or different clusters. In recent years, a number of algorithms have been proposed for enhancing clustering quality by employing such supervision. Such methods use the constraints to either modify the objective function, or to learn the distance measure. We propose a probabilistic model for semi-supervised clustering based on Hidden Mar ...

Keywords: distance metric learning, hidden Markov random fields, semi-supervised clustering

16 Recommender systems and social computing: Recommending collaboration with



 **social networks: a comparative evaluation**

David W. McDonald

April 2003 **Proceedings of the SIGCHI conference on Human factors in computing systems**

Publisher: ACM Press

Full text available: [pdf\(489.61 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Studies of information seeking and workplace collaboration often find that social relationships are a strong factor in determining who collaborates with whom. Social networks provide one means of visualizing existing and potential interaction in organizational settings. Groupware designers are using social networks to make systems more sensitive to social situations and guide users toward effective collaborations. Yet, the implications of embedding social networks in systems have not been system ...

Keywords: expertise locating, information seeking, recommender systems, social networks

17 Item-based top-*N* recommendation algorithms



 Mukund Deshpande, George Karypis

January 2004 **ACM Transactions on Information Systems (TOIS)**, Volume 22 Issue 1

Publisher: ACM Press

Full text available: [pdf\(240.61 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The explosive growth of the world-wide-web and the emergence of e-commerce has led to the development of *recommender systems*---a personalized information filtering technology used to identify a set of items that will be of interest to a certain user. User-based collaborative filtering is the most successful technology for building recommender systems to date and is extensively used in many commercial recommender systems. Unfortunately, the computational complexity of these methods grows l ...

Keywords: e-commerce, predicting user behavior, world wide web

18 Industry/government track paper: Deriving marketing intelligence from online



 discussion

Natalie Glance, Matthew Hurst, Kamal Nigam, Matthew Siegler, Robert Stockton, Takashi Tomokiyo

August 2005 **Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05**

Publisher: ACM Press

Full text available: [pdf\(629.35 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Weblogs and message boards provide online forums for discussion that record the voice of the public. Woven into this mass of discussion is a wide range of opinion and commentary about consumer products. This presents an opportunity for companies to understand and respond to the consumer by analyzing this unsolicited feedback. Given the volume, format

and content of the data, the appropriate approach to understand this data is to use large-scale web and text data mining technologies. This paper ar ...

Keywords: computational linguistics, content systems, information retrieval, machine learning, text mining

19 Distributional clustering of words for text classification 

 L. Douglas Baker, Andrew Kachites McCallum

August 1998 **Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval**

Publisher: ACM Press

Full text available:  pdf(1.07 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

20 Analyzing the effectiveness and applicability of co-training 

 Kamal Nigam, Rayid Ghani

November 2000 **Proceedings of the ninth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available:  pdf(166.63 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: co-training, expectation-maximization, learning with labeled and unlabeled data, text classification

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Search Results

BROWSE

SEARCH

IEEE XPLOR GUIDE

SUPPORT

Results for "(((newsgroup or usenet) and (weighted graph))<in>metadata)"

Your search matched 1 of 1293212 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

 e-mail printer

» Search Options

[View Session History](#)[Modify Search](#)[New Search](#) Check to search only within this results set

» Key

Display Format: Citation Citation & Abstract

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

1. A min-max cut algorithm for graph partitioning and data clustering
Ding, C.H.Q.; Xiaofeng He; Hongyuan Zha; Ming Gu; Simon, H.D.;
Data Mining, 2001. ICDM 2001, Proceedings IEEE International Conference on
29 Nov.-2 Dec. 2001 Page(s):107 - 114
Digital Object Identifier 10.1109/ICDM.2001.989507

[AbstractPlus](#) | Full Text: [PDF\(800 KB\)](#) IEEE CNF

Searching for **(newsgroup or usenet) and (weighted graph)**.

Restrict to: [Header](#) [Title](#) Order by: [Expected citations](#) [Hubs](#) [Usage](#) [Date](#) Try: [Google \(CiteSeer\)](#) [Google \(Web\)](#) [Yahoo!](#) [MSN](#) [CSB](#) [DBLP](#)

2 documents found. Order: **number of citations**.

[A Min-max Cut Algorithm for Graph Partitioning and Data..](#) - Ding, He, Zha, Gu, Simon (2001) (Correct) (1 citation)
derived. The min-max cut algorithm is tested on **newsgroup** datasets and is found to outperform other
similarities between all data objects form a **weighted graph** adjacency matrix that contains all necessary
www.nerc.gov/research/SCG/cding/papers_ps/icdm2.ps

[Analysis of Gene Expression Profiles: Class Discovery and Leaf..](#) - Ding (2002) (Correct)
method very recently developed for internet **newsgroup** clustering [11]Given a **weighted graph** G with
internet **newsgroup** clustering [11]Given a **weighted graph** G with weight matrix W, we wish to partition
[linkage.rockefeller.edu/wli/microarray/ding02.pdf](http://rockefeller.edu/wli/microarray/ding02.pdf)

Try your query at: [Google \(CiteSeer\)](#) [Google \(Web\)](#) [Yahoo!](#) [MSN](#) [CSB](#) [DBLP](#)

CiteSeer.IST - Copyright [Penn State](#) and [NEC](#)

The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Web

Results 1 - 5 of about 6 for **newsgroup and cross-post and "weighted graph"**. (0.38 seconds)

Tip: Try removing quotes from your search to get more results.

[\[PDF\] Exploring the Community Structure of Newsgroups](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

content of the **newsgroups** to which they are posted. We define the **cross-post** graph as a **weighted graph** with vertices representing the **newsgroups** and ... research.microsoft.com/~jchayes/Papers/cluster.pdf - Similar pages

[United States Patent Application: 0040267686](#)

the **weighted graph** represents A subset of the **newsgroups** as vertices of the graph.

... [0012] A **weighted graph** is created based upon the **newsgroups** ...

appft1.uspto.gov/netacgi/nph-Parser?Sect1=PTO2& Sect2=HITOFF&p=38&u=%2Fnetacgi%2FPTO%2Fsearch... - 102k - Supplemental Result - Cached - Similar pages

[www.grahamkendall.net/Math/Math%20Newsgroups/mm-11...](#)

291k - Supplemental Result - Cached - Similar pages

[lists.humbug.org.au/archives/chat/1997-October.txt.gz](#)

1100k - Supplemental Result - Cached - Similar pages

[Sam Sloan's Chess Lesson Video Posted on the Internet](#)

Well, especially as most people will be subscribed to all the **rec.games.Chess**

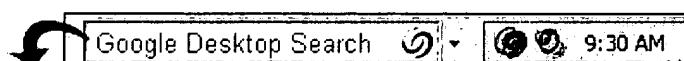
newsgroups. However, people talking about it are presumably what Sam wants, ...

www.avlerchess.com/chess-analysis/Sam_Sloans_Chess_Lesson_Video_Posted_on_the_Internet_74802.html - 245k - Supplemental Result - Cached - Similar pages

In order to show you the most relevant results, we have omitted some entries very similar to the 5 already displayed.

If you like, you can [repeat the search with the omitted results included](#).

Try your search again on [Google Book Search](#)



Free! Instantly find your email, files, media and web history. [Download now](#).

The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Tip: Try removing quotes from your search to get more results.

Your search - **newsgroup and cross-post and "weighted graph"** - did not match any articles.

Suggestions:

- Make sure all words are spelled correctly.
- Try different keywords.
- Try more general keywords.
- Try fewer keywords.
- [Try your query on the entire web.](#)

[Google Home](#) - [About Google](#) - [About Google Scholar](#)

©2005 Google

The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Tip: Try removing quotes from your search to get more results.

Your search - **(newsgroup OR usenet) AND cross-post and "weighted graph"** - did not match any articles.

Suggestions:

- Make sure all words are spelled correctly.
- Try different keywords.
- Try more general keywords.
- Try fewer keywords.
- [Try your query on the entire web.](#)

[Google Home](#) - [About Google](#) - [About Google Scholar](#)

©2005 Google

Terms used [newsgroup](#) and [cross post](#) and [weighted graph](#)

Found 40 of 169,166

Sort results
by relevance Save results to a Binder[Try an Advanced Search](#)Display
results expanded form Search Tips
 Open results in a new window[Try this search in The ACM Guide](#)

Results 1 - 20 of 40

Result page: [1](#) [2](#) [3](#) [next](#)

Relevance scale

1 [Industry/government track posters: Exploring the community structure of newsgroups](#)

Christian Borgs, Jennifer Chayes, Mohammad Mahdian, Amin Saberi

August 2004 **Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '04****Publisher:** ACM PressFull text available: [pdf\(3.11 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We propose to use the community structure of Usenet for organizing and retrieving the information stored in newsgroups. In particular, we study the network formed by cross-posts, messages that are posted to two or more newsgroups simultaneously. We present what is, to our knowledge, by far the most detailed data that has been collected on Usenet cross-postings. We analyze this network to show that it is a small-world network with significant clustering. We also present a spectral algorithm which ...

Keywords: clustering, spectral method, usenet**2** [The dynamics of mass interaction](#)

Steve Whittaker, Loren Terveen, Will Hill, Lynn Cherny

November 1998 **Proceedings of the 1998 ACM conference on Computer supported cooperative work****Publisher:** ACM PressFull text available: [pdf\(913.66 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**Keywords:** FAQS, Usenet, common ground, conversation, empirical, mass interaction, moderation, netiquette, newsgroups**3** [Collaborative Filtering: Observed behavior and perceived value of authors in usenet newsgroups: bridging the gap](#)

Andrew T. Fiore, Scott Lee Tiernan, Marc A. Smith

April 2002 **Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves****Publisher:** ACM PressFull text available: [pdf\(383.54 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we describe an evaluation of behavioral descriptors generated from an analysis of a large collection of Usenet newsgroup messages. The metrics describe aspects of newsgroup authors' behavior over time; such information can aid in filtering, sorting, and recommending content from public discussion spaces like newsgroups. To assess the value of a variety of these behavioral descriptors, we compared 22 participants' subjective evaluations of authors whose messages they read to behavio ...

Keywords: behavioral indicators, discussions, persistent conversations, social accounting, social cyberspaces

4 Usenet News for electronic information sharing

 Glenn Leavell
June 1995 **ACM SIGUCCS Newsletter**, Volume 25 Issue 1-2

Publisher: ACM Press

Full text available:  pdf(595.24 KB) Additional Information: [full citation](#), [abstract](#)

UCNS is pleased to announce the availability of a new Usenet News server on campus, news.uga.edu. Usenet News is an electronic public forum in which you can participate in discussions and exchange information with millions of people around the world. This article will explain what Usenet News is, how it works, and how you can take advantage of this tremendous resource.

5 Supporting community and building social capital: Tools for navigating large social

 cyberspaces

Marc Smith
April 2002 **Communications of the ACM**, Volume 45 Issue 4

Publisher: ACM Press

Full text available:  pdf(256.18 KB)
 html(21.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The Netscan project helps online participants form cooperative relationships by offering a better sense of the other players involved.

6 GroupLens: an open architecture for collaborative filtering of netnews

 Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, John Riedl
October 1994 **Proceedings of the 1994 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available:  pdf(1.32 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Collaborative filters help people make choices based on the opinions of other people. GroupLens is a system for collaborative filtering of netnews, to help people find articles they will like in the huge stream of available articles. News reader clients display predicted scores and make it easy for users to rate articles after they read them. Rating servers, called Better Bit Bureaus, gather and disseminate the ratings. The rating servers predict scores based on the heuristic that people wh ...

Keywords: Usenet, collaborative filtering, electronic bulletin boards, information filtering, netnews, selective dissemination of information, social filtering, user model

7 GroupLens: applying collaborative filtering to Usenet news

 Joseph A. Konstan, Bradley N. Miller, David Maltz, Jonathan L. Herlocker, Lee R. Gordon, John Riedl
March 1997 **Communications of the ACM**, Volume 40 Issue 3

Publisher: ACM Press

Full text available:  pdf(343.16 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

8 Building task-specific interfaces to high volume conversational data

 Loren G. Terveen, William C. Hill, Brian Amento, David McDonald, Josh Creter
March 1997 **Proceedings of the SIGCHI conference on Human factors in computing systems**

Publisher: ACM Press

Full text available:  pdf(908.00 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: Netnews, Usenet, World Wide Web, collaborative filtering, computer-supported cooperative work, data mining, human interface, human-computer interaction, organizational computing, resource discovery, social filtering

9 Text Extraction and Summarization: Text classification in a hierarchical mixture model for small training sets

 Kristina Toutanova, Francine Chen, Kris Popat, Thomas Hofmann
October 2001 **Proceedings of the tenth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available:  [pdf\(1.40 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Documents are commonly categorized into hierarchies of topics, such as the ones maintained by Yahoo! and the Open Directory project, in order to facilitate browsing and other interactive forms of information retrieval. In addition, topic hierarchies can be utilized to overcome the sparseness problem in text categorization with a large number of categories, which is the main focus of this paper. This paper presents a *hierarchical mixture model* which extends the standard naive Bayes classif ...

10 An adaptive k -nearest neighbor text categorization strategy

 Li Baoli, Lu Qin, Yu Shiwen
December 2004 **ACM Transactions on Asian Language Information Processing (TALIP)**, Volume 3 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(285.94 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

k is the most important parameter in a text categorization system based on the k -nearest neighbor algorithm (k NN). To classify a new document, the k -nearest documents in the training set are determined first. The prediction of categories for this document can then be made according to the category distribution among the k nearest neighbors. Generally speaking, the class distribution in a training set is not even; some classes may have more samples than others. ...

Keywords: k -nearest neighbor algorithm, machine learning, text categorization, text classification

11 Involving remote users in continuous design of web content

 William C. Hill, Loren G. Terveen
August 1997 **Proceedings of the conference on Designing interactive systems: processes, practices, methods, and techniques**

Publisher: ACM Press

Full text available:  [pdf\(915.47 KB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: Usenet, World Wide Web, collaborative filtering, computer-supported cooperative work, end user modification, human interface, human-computer interaction, organizational computing, participatory design, remote evaluation, resource discovery, social filtering

12 PHOAKS: a system for sharing recommendations

 Loren Terveen, Will Hill, Brian Amento, David McDonald, Josh Creter
March 1997 **Communications of the ACM**, Volume 40 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(328.21 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 **Research session 5: data mining / transaction management: A divide-and-merge methodology for clustering**

David Cheng, Santosh Vempala, Ravi Kannan, Grant Wang

June 2005 **Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available: [pdf\(791.76 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

We present a divide-and-merge methodology for clustering a set of objects that combines a top-down "divide" phase with a bottom-up "merge" phase. In contrast, previous algorithms either use top-down or bottom-up methods to construct a hierarchical clustering or produce a flat clustering using local search (e.g., k -means). Our divide phase produces a tree whose leaves are the elements of the set. For this phase, we use an efficient spectral algorithm. The merge phase quickly finds an optim ...

14 **All ways aware: Stimulating social engagement in a community network**

David R. Millen, John F. Patterson

November 2002 **Proceedings of the 2002 ACM conference on Computer supported cooperative work**

Publisher: ACM Press

Full text available: [pdf\(384.26 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

One of the most challenging problems facing builders and facilitators of community networks is to create and sustain social engagement among members. In this paper, we investigate the drivers of social engagement in a community network through the analysis of three data sources: activity logs, a member survey, and the content analysis of the conversation archives. We describe three important ways to encourage and support social engagement in online communities: through system design elements suc ...

Keywords: community networks, community-support systems, empirical studies, social awareness, social engagement

15 **Analyzing the effectiveness and applicability of co-training**

Kamal Nigam, Rayid Ghani

November 2000 **Proceedings of the ninth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: [pdf\(166.63 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: co-training, expectation-maximization, learning with labeled and unlabeled data, text classification

16 **Research track: Information-theoretic co-clustering**

Inderjit S. Dhillon, Subramanyam Mallela, Dharmendra S. Modha

August 2003 **Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining**

Publisher: ACM Press

Full text available: [pdf\(1.96 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Two-dimensional contingency or co-occurrence tables arise frequently in important applications such as text, web-log and market-basket data analysis. A basic problem in contingency table analysis is *co-clustering: simultaneous clustering* of the rows and columns. A novel theoretical formulation views the contingency table as an empirical joint probability distribution of two discrete random variables and poses the co-clustering problem as an optimization problem in *information theory*

Keywords: co-clustering, information theory, mutual information

17 [Topic search for intelligent network news reader HISHO](#) 
Hiromi Ozaku, Kiyotaka Uchimoto, Masaki Murata, Hitoshi Ishara
March 2000 **Proceedings of the 2000 ACM symposium on Applied computing - Volume 1**
Publisher: ACM Press
Full text available:  [pdf\(989.75 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

18 [Research track poster: Co-clustering by block value decomposition](#) 
Bo Long, Zhongfei (Mark) Zhang, Philip S. Yu
August 2005 **Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05**
Publisher: ACM Press
Full text available:  [pdf\(775.70 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Dyadic data matrices, such as co-occurrence matrix, rating matrix, and proximity matrix, arise frequently in various important applications. A fundamental problem in dyadic data analysis is to find the hidden block structure of the data matrix. In this paper, we present a new co-clustering framework, block value decomposition(BVD), for dyadic data, which factorizes the dyadic data matrix into three components, the row-coefficient matrix **R**, the block value matrix **B**, and the column-c ...

Keywords: block value decomposition (BVD), clustering, co-clustering, dyadic data, hidden block structure, matrix decomposition, non-negative block value decomposition (NBVD)

19 [Research track poster: Model-based overlapping clustering](#) 
Arindam Banerjee, Chase Krumpelman, Joydeep Ghosh, Sugato Basu, Raymond J. Mooney
August 2005 **Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05**
Publisher: ACM Press
Full text available:  [pdf\(837.60 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

While the vast majority of clustering algorithms are partitional, many real world datasets have inherently overlapping clusters. Several approaches to finding overlapping clusters have come from work on analysis of biological datasets. In this paper, we interpret an overlapping clustering model proposed by Segal et al. [23] as a generalization of Gaussian mixture models, and we extend it to an overlapping clustering model based on mixtures of any regular exponential family distribution and the c ...

Keywords: Bregman divergences, exponential model, graphical model, high-dimensional clustering, overlapping clustering

20 [Index structures for selective dissemination of information under the Boolean model](#) 
Tak W. Yan, Héctor García-Molina
June 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 2
Publisher: ACM Press
Full text available:  [pdf\(2.03 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The number, size, and user population of bibliographic and full-text document databases are rapidly growing. With a high document arrival rate, it becomes essential for users of such databases to have access to the very latest documents; yet the high document arrival rate also makes it difficult for users to keep themselves updated. It is desirable to allow users to submit profiles, i.e., queries that are constantly evaluated, so that they will be automatically informed of new additions tha ...

Useful downloads: [!\[\]\(efee95c0f3d6f57aebda4c234cf0ca9b_img.jpg\) Adobe Acrobat](#) [!\[\]\(870b91d7ec09688d7e47f3c7817fdf66_img.jpg\) QuickTime](#) [!\[\]\(4fb4a1c91ad85efe4ccc7cf981ce94c0_img.jpg\) Windows Media Player](#) [!\[\]\(ce63ca28962cc2e105c0c64de024bc81_img.jpg\) Real Player](#)